EXHIBIT "A"

Disclosure AUS8-2000-1323

Created By: Rabindranath Dutta **Created On:** [redacted]
Last Modified By: Rabindranath Dutta **Last Modified On:** [redacted]

*** IBM Confidential ***

Required fields are marked with the asterisk (*) and must be filled in to complete the form .

Summary

Status	Under Evaluation
Processing Location	AUS
Functional Area	58 - TIVOLI (B. Yellin, David Murphy, P. Morowski, L. Orecklin, Watkins, D.Kerr)
Attorney/Patent Professional	Jeff LaBaw/Austin/IBM
IDT Team	John Sweitzer/Tivoli Systems; Sebastian Hassinger/Tivoli Systems; Nicole Harbour/Tivoli Systems
Submitted Date	[redacted]
Owning Division Select	TIV
PVT Score Calculate	To calculate a PVT score, use the 'Calculate PVT' button.
Incentive Program	
Lab	
Technology Code	

Inventors with Lotus Notes IDs

Inventors: Rabindranath Dutta/Tivoli Systems, Janani Janakiraman/Tivoli Systems

Inventor Name > denotes primary contact	Inventor Serial	Div/Dept	Manager Serial	Manager Name
Dutta, Rabindranath	978208	46/LYKS	388393	Zhang, R.P. (Randy)
> Janakiraman, Janani	0A3131	46/AHCS	998634	Willard, R.G. (Ron)

Inventors without Lotus Notes IDs

IDT Selection

IDT Team:	Attorney/Patent Professional:
John Sweitzer/Tivoli Systems Sebastian Hassinger/Tivoli Systems Nicole Harbour/Tivoli Systems	Jeff LaBaw/Austin/IBM

Response Due to IP&L : 11/08/2000

Main Idea

Title of disclosure (in English)

Method to create a rating server of Gnutella nodes based on usage and download patterns.

Idea of disclosure

1. Describe your invention, stating the problem solved (if appropriate), and indicating the advantages of using the invention.

Background

Very recently the traditional Web server - Web browser paradigm has been challenged by distributed file systems based solutions like Gnutella that provide means for **peer-to-peer exchange of data.** Essentially from the client-server model we are progressing towards a a fully-distributed information-sharing technology. A functional description of Gnutella networks as used for search is given below:

"When you're looking for something on the Internet, you generally ask a search engine, such as AltaVista or Yahoo, to find it for you. The engine checks the Web sites it knows about (the average search engine actually searches less than 20 percent of all the sites on the Internet). Users of peer-to-peer software such as Gnutella in essence form a search engine of their own that expands its search exponentially. When a Gnutella user has a query, the software sends it to 10 computers on the network. If the first 10 computers don't have the file, each computer sends it to 10 other computers and so on until, designers say, an estimated million computers would be looking for it in just five to 10 seconds. The program could theoretically check every site on the Web." (excerpt from a CNET news article).

Essentially in peer-to-peer Gnutella networks, a client enters the IP addresses of several other clients and they communicate directly without a central server. The connection is extended very rapidly in depth by each client connecting to other clients as explained above.

Peer to peer networking for information is likely to be a revolutionary area in the Internet. To quote from the article on the subject reproduced in the Appendix: "... the potential for peer-to-peer systems is just starting to be realized. He also said that Napster and its cousins, Gnutella and Freenet, were comparable to Mosaic, the first Web browser from the University of Illinois that made the Web accessible to the masses for the first time. "

For official details on Gnutella and peer-to-peer networking see the site: http://gnutella.wego.com. *While we are writing this invention with reference to Gnutella the attorney writing this application should write it with reference to generalized peer-to-peer information sharing networks..* The reference to Gnutella is just to aid comprehension.

This disclosure is a series of disclosures on Gnutella like "peer-to-peer" data exchange. See Appendix for a description of the other disclosures.

Problem

In a distributed file sharing network like Gnutella, each node participating in the search acts as a client as well as a server. In the server mode, the node exports files that it wants to share with other users. In the client mode, the nodes perform a search for a file. The nodes maintain a list of TCP/IP addresses (i.e. starting connection points) and when performing a search send the search request to the nodes in their connection list. If none of these nodes have the requested file, these nodes in turn send the request to the set of connection addressed maintained by each of them.

There is no way to determine the starting set of nodes which are worth connecting to in the context of the search item. The search might fan out to nodes whose content is in no way related to the item being searched for. Currently there is no clear method by which central servers can create such starting points although ad hoc mechanism exist in the case of some sites that provide starting Gnutella nodes (for example www.gnutellahosts.com). Such prior art solutions do not categorize content (which is of course part of another invention that we have submitted).

Solution

To solve the above mentioned problem, our invention proposes the establishment of a rating mechanism for all the nodes participating in the search. Client nodes interested in using the rating mechanism, register with centralized rating servers (although interaction with a central server was not something that Gnutella intended we believe from our experience in using Gnutella that such central servers will be required in some form for enhancing Gnutella. Since

AUS8-2000-1323 Method to create a rating server of Gnutella nodes based on usage and download patterns. - continued

Gnutella started in the aftermath of lawsuits with Napster the technology has gravitated towards attempting to avoid any kind of centralized servers. However in general peer-to-peer computing with a limited amount of centralized actives is a superior paradigm that is easy to architect). The rating servers maintain a list of starting connection points categorized by the keyword in the search string. This database is built by prior experience of the clients optionally registered with this rating server. Before starting a search, the client can consult the rating server to get an Initial set of starting points for the search based on the keyword in their search string. The populating of the Gnutella nodes with the appropriate TCP/IP addresses for beginning the search is automatic from the server. For example the rating server has "keywords" and corresponding "IP addresses of good candidates of starting a Gnutella search"

In the invention we focus on the methods to build the database. Other inventions in the series are on aspects other than building the database in the rating server.

Rating methodology:

The client nodes retrieve files from other nodes based on the search string. The rating software examines and monitors the results of the searches performed by the client and checks to see if the search result files are downloaded, examined by user and how many time they are played or used by the user. This information is stored as a keyword/file mapping for the machine it has been downloaded from. If the file is used repeatedly by the user, the ranking of the file for the keyword is increased.

The client node, returns this rating information at regular intervals to the rating server which maintains the master DB of this connection rating information.

All the clients registered with the rating server can make use of this rating information in gathering a set of initial nodes to connect to.

See figure below for the flow of rating information between the client and server.

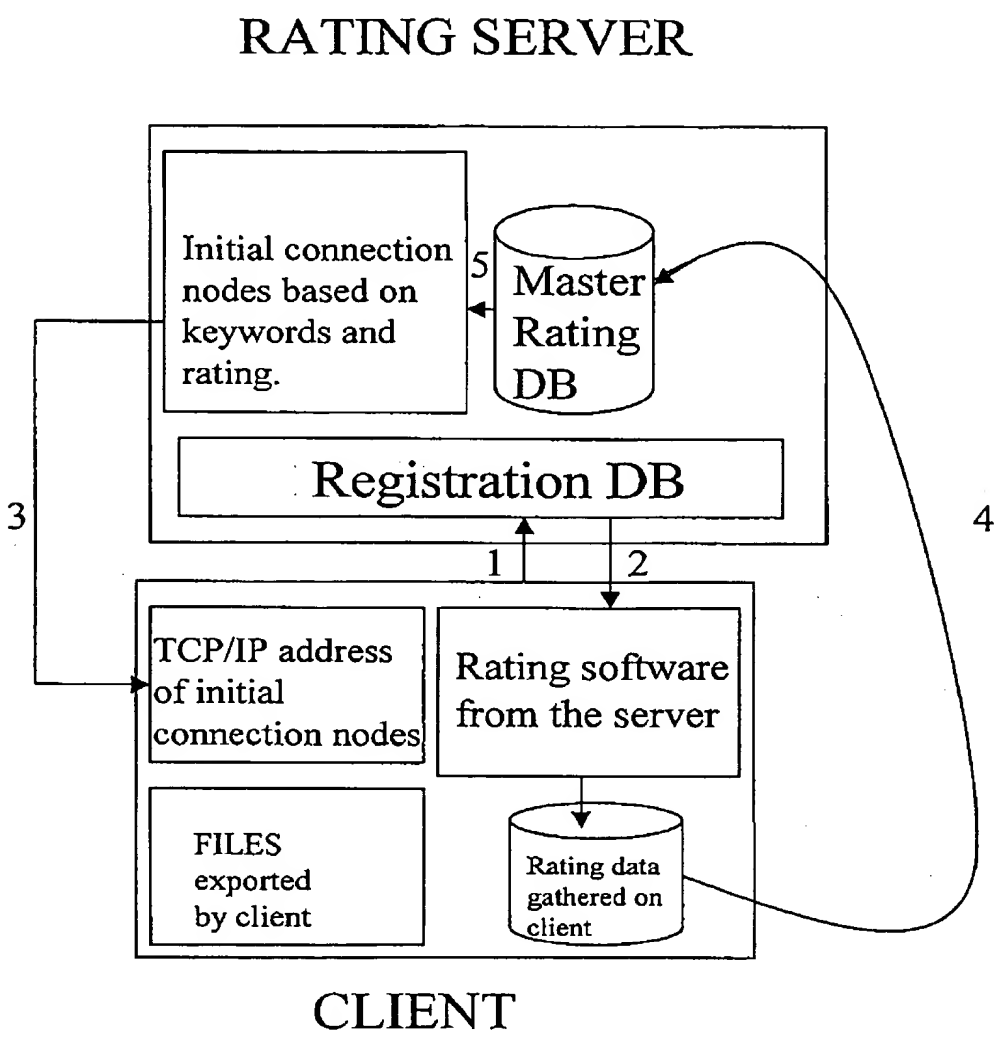

* Client would have client interest files/areas set up.
* Server has a list of initial connection hosts categorized by keyword.

1. Client registers with the rating server
2. Rating server installs a rating software on the client machine. This software is the one which analyzes the gnutella search result traffic coming into the client and watches to see how much it is used.
3. Before starting a search, the client downloads the set of connection points based on the keyword in its search string.
4. Client periodically sends the rating information gathered back to the rating server.
5. The rating server collects the rating information from all its registered clients and updates its rating database.

Other clients who subscribe to this methodology can also use these services.

To illustrate with an example:

- Client wants to search for "Music by Beethoven"
- Client registers with a rating server and downloads a set of initial TCP/IP addresses to start its search.
- When the client registers with the rating server, the rating server also puts some rating software on the client machine which monitors the search result files based on usage.
- The information gathered by the client is returned to the server , for instance if the user plays a certain music file many times, it gets a higher rating. If he downloads a file but never plays it , it will have a lower rating.
- Other clients looking for similar music files, can make use of the prior experience of clients registered with the rating server to get a set of initial TCP/IP addresses to start their search.

Advantages

- Business Method: Rating of gnutella nodes based on usage
- The concept of rating can be extended to regular search engines.

Claims

(broadest claim) Method to create a rating server for Gnutella nodes based on usage and download patterns.

(The method is different from that submitted in disclosure 1 shown in the appendix. Disclosure 1 focusses on
a) Providing a list of accesible sites based on the category of the file
b) Claiming the GUI aspects as shown in one of the figures

The current disclosure focusses on how the rating server actually creating the ratings based on download and usage patterns.)

Appendix:

Here is the text of the other disclosures submitted in this arena. Disclosure 2 is going for filing and Disclosure 1 is still under review.

DISCLOSURE 1 (disclosure-946)
Augmenting prior-art search engine results with Gnutella client addresses

PROBLEM:

Since the Web is a dynamic environment where content is often being added, updated and moved, it is very difficult for search engine databases to be kept up to date. Statistics show that at most 30% of the web pages of the Web universe is stored in the search engine databases. ***We need a mechanism to be able to search the other 70% of the published content on the web. But since information is so distributed , we need a way to be able to search the information space in an organized fashion.***

SOLUTION:
Our solution provides a way for the prior art search engines to augment the conventional search engines with Gnutella search mechanisms. For example, search engines like Altavista, would not only return the static hyperlinked pages from their databases, but they would also provide a starting point to begin the Gnutella search of the rest of the information on client computers spread all over the Internet. Web

Portals like Yahoo, will be able to categorize the information on the individual nodes based on shared interest and would provide starting points (Gnutellahosts) for performing a Gnutella search.

If a bunch of hosts (e.g. see figure below) are connected, the network space can be carved out into sub-networks (see ovals) and the portal can provide hostnames of the major root node (s of these sub networks e.g. Node "a" and Node "x" in our figure. The nodes of these networks in the figure below are client computers with information to share. They are linked by peer-to-peer networking. For example the client d when it searches, first accesses b and then b accesses a, c, k and g and so on. In the figure given below the Gnutella network is partitioned by our invention server into sub-networks of different area. Such partitioning is the subject of a different invention that we will be submitting and is not central to the current invention. The current invention proceeds on the assumption that there are methods available to generally classify into broad groups the clients located in the Gnutella network. Essentially the simplest solution is to have a group of reference clients that classifies information located on clients by connecting. While dynamic IPs may provide some problems, Gnutella could be augmented to provide more information identifying each client that its dynamic IP address only. Currently, the dynamic IP is used to preserve some anonymity - in particular because of Napster related copyright lawsuits etc. However, we believe that once Gnutella networks become popular for information that is not subject to copyright disputes the need for anonymity would get reduced.

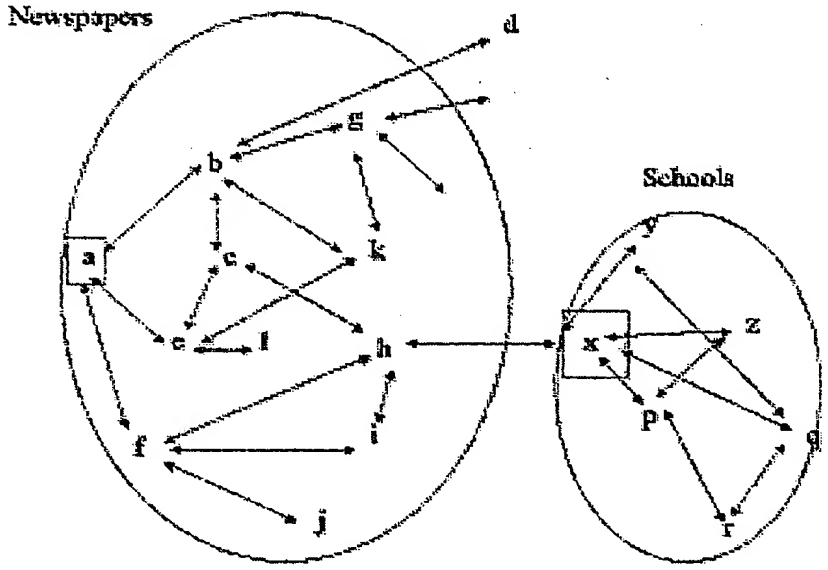

Below is an example of the page presented by search engines by our invention which incorporates the gnutella search along with the conventional prior art search.

In the figure below when a user clicks on "Gnutella IPs at current time" or to the "Dynamic Gnutella search", immediately the user's Gnutella client will connect to these IPs for a more extended search. The link will also restore the user's "Gnutella" client window to the foreground so that he can continue the search in the Gnutella space.

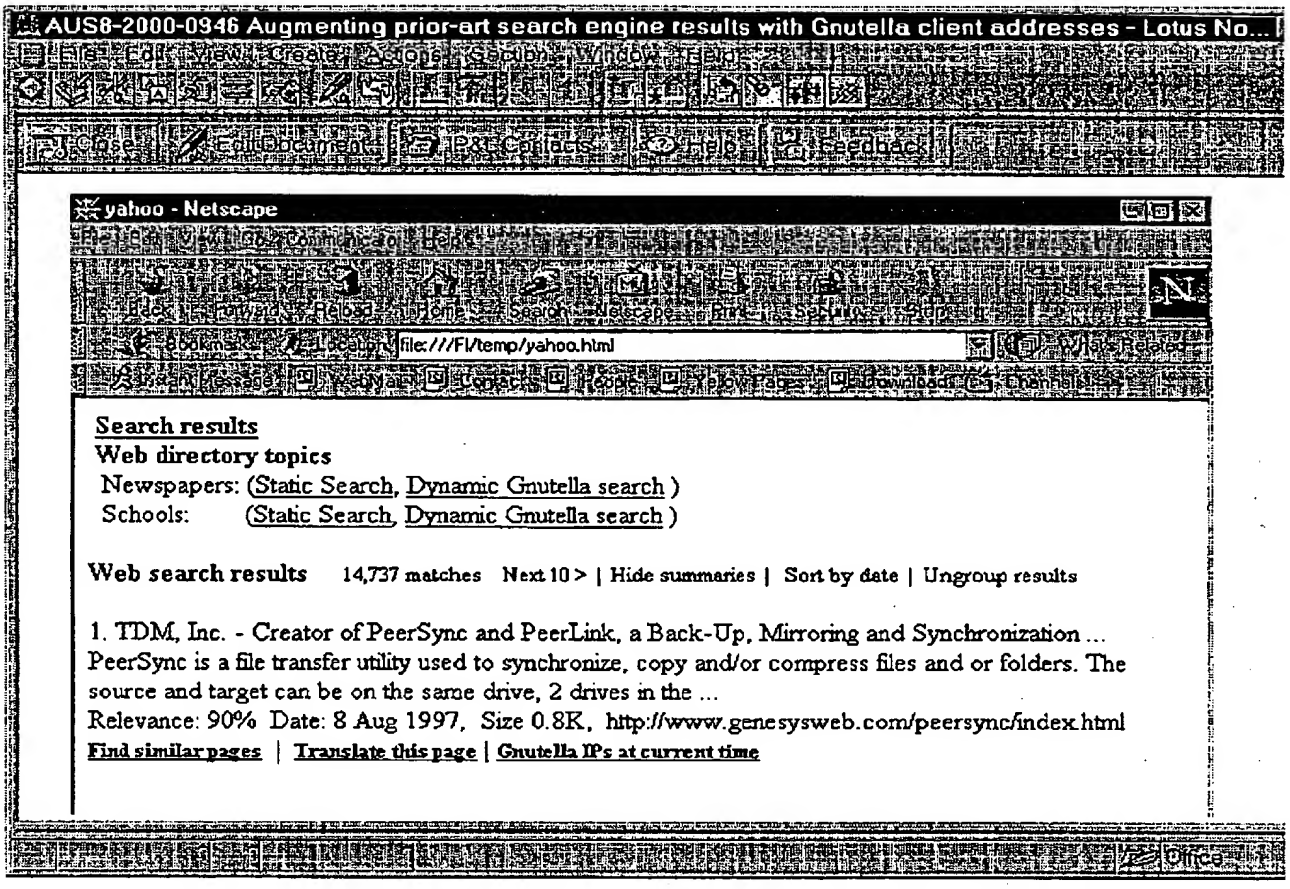

In essence, the user is first searching the limited amount of Internet information indexed by search engines and then performing a more extended search using the client to search the rest of the information content on the Internet that may be present in other Gnutella clients.

The advantages of the invention is that the load on search engines are reduced and that the user can search an extended universe of Internet data. It is different from prior art systems such as Napster which merely provided links to the clients on which information existed.

CLAIMS

(broadest claims)

- A method to augment conventional search engine results with addresses of representative nodes of peer-to-peer network based file sharing systems so as to increase the amount of information available to users.

- The business method where Gnutella client "root" nodes of the subnetwork will receive incentives, money etc. for participating with the search engine.

Note - The rush to patent Gnutella based technologies is likely to accelerate significantly as major corporations get involved in the effort once it attains critical mass in the user community. The success of Napster and Gnutella in college dorms is an indicator of the future trend of the Internet. We have written up a sequence of about 20 inventions in this area covering many aspects of the technology and will be submitting many of them in the near future to the evaluation board. The

ones that are being submitted next relate to the following:
1. dynamic reordering of Gnutella nodes for load balancing
2. thumbnail display of client characteristics and query of clients
3. categorization of Gnutella subnetworks based on content etc.
4. queuing of Gnutella clients with authorization tickets.

Appendix

(http://www.wired.com/news/print/[redacted].html)

Intel Says: Think Like Napster
by Leander Kahney

[redacted]
SAN JOSE, Calif. -- An Intel executive on Thursday said Napster and programs like it will transform the
Internet as profoundly as the first Web browser.

Patrick Gelsinger, chief technology officer of Intel's architecture group, said Napster and other
peer-to-peer networking technologies are "a revolution that will change computing as we know it."

During a keynote speech at the Intel Developer Forum here, Gelsinger called on the computing industry to
cooperate in hammering out standards for peer-to-peer networking, which he said will usher in the "next
era" of computing.

Peer-to-peer networks allow any Net-connected computer to talk to any other without having to go
through centralized servers.

In Napster's case, a peer-to-peer network is used to share music files between users. The music isn't
stored on central servers, but on the individual user's machines.

In a dig at Napster, currently embroiled in a legal battle with the recording industry, Gelsinger said Intel
paid for the music clip that was playing as he walked onstage.

As well as sharing files -- music, pictures or video -- peer-to-peer systems can also be used to farm out
chunks of complex problems between thousands of machines. The SETI@home screen saver, for example,
analyzes radio signals from space for signs of intelligent life.

**Gelsiger said the potential for peer-to-peer systems is just starting to be realized. He also said that
Napster and its cousins, Gnutella and Freenet, were comparable to Mosaic, the first Web browser from
the University of Illinois that made the Web accessible to the masses for the first time.**

Before Mosaic there were 50 Web servers, Gersinger said. A year after it appeared, there were 10,000.

"Peer-to-peer could usher in the next era of the Internet much the same way Mosaic ushered in the last,"
he said.

Looking forward, Gelsinger said that just as it was impossible to envision Amazon and eBay when Mosaic
was created, it is impossible to foresee all the applications for peer-to-peer computing.

He did, however, suggest a few -- collaboration, file sharing, intelligent agents like virus sentinels, and
distributed computing.

"We think the applications are nearly limitless," he said.

Gelsinger demonstrated Intel's NetBatch, a distributed computing system that uses spare cycles on the company's servers and workstations to solve chip design problems.

He said an internal study showed that company workstations were idle 75 percent of the time, and servers were idle 50 percent of the time. In 1990, Intel started NetBatch with a few hundred systems. It has now grown to 10,000 computers. NetBatch utilizes 80 percent of the company's total computing power, solves 2.7 million problems a month, and has saved half a billion dollars in equipment costs.

"We are consistently accelerating our chip schedules because of the mammoth, mammoth, computing power we can apply," he said.

The system has also led to big savings in storage requirements and better network performance, he added.

Gelsinger announced the formation of the peer-to-peer working group www.peer-to-peerwg.org, a consortium of 19 companies that are mostly distributed computing startups but also including IBM and Hewlett Packard. The group intends to establish standards for management and security, define protocols, and hammer out solutions.

Gelsinger said Microsoft was talking about joining.

Nathan Brookwood, an analyst with Insight 64, said that Intel, the world's largest chipmaker, clearly sees peer-to-peer networking as a way of driving demand for faster processors.

"Most people are happy with a 500 or 600 MHz machine, which is fast enough for most tasks," he said. "But Intel needs to make these machines seem slow. How do they do that? New workloads and applications. ... It's a way of forcing people to upgrade their machines."

DISCLOSURE 2 (disclosure 979)
Thumbnail display of content in Gnutella clients and consequent improvement in reducing fanout factor for searching
Background
Very recently the traditional Web server - Web browser paradigm has been challenged by distributed file systems based solutions like Gnutella that provide means for **peer-to-peer exchange of data.** Essentially from the client-server model we are progressing towards a a fully-distributed information-sharing technology. A functional description of Gnutella networks as used for search is given below:

"When you're looking for something on the Internet, you generally ask a search engine, such as AltaVista or Yahoo, to find it for you. The engine checks the Web sites it knows about (the average search engine actually searches less than 20 percent of all the sites on the Internet). Users of peer-to-peer software such as Gnutella in essence form a search engine of their own that expands its search exponentially. When a Gnutella user has a query, the software sends it to 10 computers on the network. If the first 10 computers don't have the file, each computer sends it to 10 other computers and so on until, designers say, an estimated million computers would be looking for it in just five to 10 seconds. The program could theoretically check every site on the Web." (excerpt from a CNET news article).

Essentially in peer-to-peer Gnutella networks, a client enters the IP addresses of several other clients and they communicate directly without a central server. The connection is extended very rapidly in depth by each client connecting to other clients as explained above.

Additional details of Gnutella can be obtained from the disclosure "Augmenting prior-art search engine results with Gnutella client addresses."

Problem
In a peer-to-peer datasharing network, when a connection is established with a peer, there is no way for the user to determine quickly if it is worth browsing the content on the peer. Since the search fans out to a widely distributed network , the search can often take a client to nodes that do not share the same interest as the user. Also the fan out of nodes is limited and eliminating unproductive connections will speed up the search. Our solution attempts to solve the above mentioned problems.

Solution
Having a thumbnail sketch of the peer's characteristics will help the client decide if it is of any use to pursue browsing the content on the peer or the nodes connected to the peer. The solution is explained in somewhat greater detail below.

For peer-to-peer data sharing networks example Gnutella, as soon as Gnutella client A establishes a connection to Gnutella Client B, B should immediately in one of the early packets display to A its characteristics. The characteristics would include but in not limited to:
1. when is a typical good time to connect
2. thumbnail sketch of information area within N links deep.
3. thumbnail characteristics of what areas it covers
4. speed, fan-out/fan-in load of the network connectivity it can support

We believe that if the patent covered the above four characteristics it would be adequate to prevent others from arriving at an alternative solution without violating at least some of the claims of the patent. Essentially the above characteristics are the main pieces of information needed for searching and restricting fanout. There may be other characteristics that are less important.

The information can be displayed in a pop up box which appears when the user hovers over the connections established with the peers. This provides a mechanism by which the user can quickly disconnect a connection if he feels that the content on the site it is connected to is not of interest. Claiming the GUI interface would make it difficult for others to arrive at alternative methods of solving the problem as the GUI information can only be represented effectively in a relatively small number of ways.

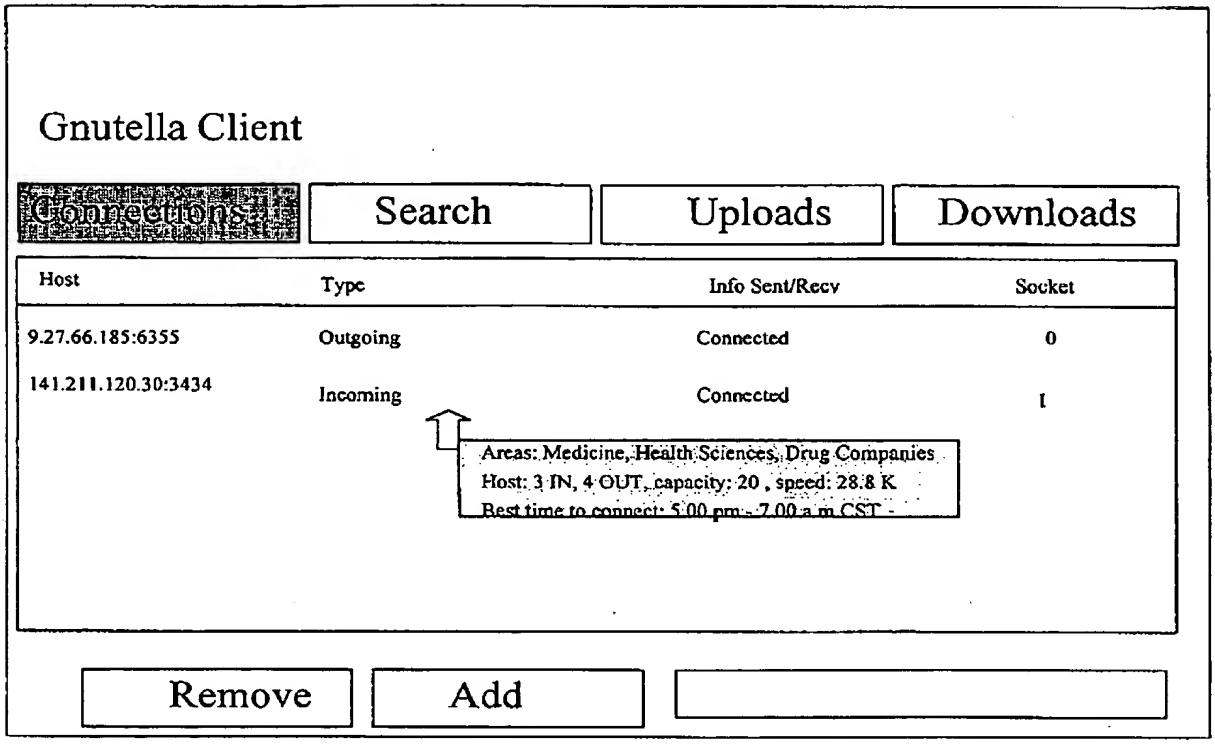

CLAIMS

1. Method to display peer client characteristics in peer-to-peer content sharing networks on the Internet.

(dependent claims)
where the client characteristics are
A. when is a typical good time to connect
B. thumbnail sketch of information area within N links deep.
C. thumbnail characteristics of what areas it covers
D. speed, fan-out/fan-in load of the network connectivity it can support

2. How does the invention solve the problem or achieve an advantage,(a description of "the invention", including figures inline as appropriate)?

3. If the same advantage or problem has been identified by others (inside/outside IBM), how have those others solved it and does your solution differ and why is it better?

4. If the invention is implemented in a product or prototype, include technical details, purpose, disclosure details to others and the date of that implementation.

AUS8-2000-1323 Method to create a rating server of Gnutella nodes based on usage and download patterns. - continued

***Critical Questions (Questions 1 - 7 must be answered)**

***Question 1** On what date was the invention workable? [redacted] format the date as MM/DD/YYYY. (Workable means i.e. when you know that your design will solve the problem)	
***Question 2** Is there any planned or actual publication or disclosure of your invention to anyone outside IBM?	○ Yes ● No
If yes, Enter the name of each publication or patent and the date published below. Publication/Patent: Date Published or Issued:	
Are you aware of any publications, products or patents that relate to this invention?	○ Yes ● No
If yes, Enter the name of each publication or patent and the date published below. Publication/Patent: Date Published or Issued:	
***Question 3** Has the subject matter of the invention or a product incorporating the invention been sold, used internally in manufacturing, announced for sale, or included in a proposal?	○ Yes ● No
Is a sale, use in manufacturing, product announcement, or proposal planned?	○ Yes ● No
If Yes, identify the product if known and indicate the date or planned date of sale, announcements, or propposal and to whom the sale, announcement or proposal has been or will be made. Product: Version/Release: Code Name: Date: To Whom: If more than one, use cut and paste and append as necessary in the field provided.	
***Question 4** Was the subject matter of your invention or a product incorporating your invention used in public e.g. outside IBM or in the presence of non-IBMers?	○ Yes ● No
If yes, give a date. Please format the date as MM/DD/YYYY	
***Question 5** Have you ever discussed your invention with others not employed at IBM?	○ Yes ● No
If yes, identify individuals and date discussed. Fill in the text area with the following information: the names of the individuals, the employer, date discussed, under CDA, and CDA #	
***Question 6** Was the invention, in any way, started or developed under a government contract or project?	○ Yes ● No ○ Not sure
If Yes, enter the contract number	
***Question 7** Was the invention made in the course of any alliance, joint development or other contract activities?	○ Yes ● No ○ Not Sure

AUS8-2000-1323 Method to create a rating server of Gnutella nodes based on usage and download patterns. - continued

If Yes, enter the following: Name of Alliance, Contractor or Joint Developer
Contract ID number
Relationship contact name
Relationship contact E-mail
Relationship contact phone

Question 8 Have you submitted, or are you aware of, any related disclosure submission?	○ Yes ● No
If Yes, please provide the title and docket or disclosure number below:	

Question 9
What type of companies do you expect to compete with inventions of this type? *Check all that apply.*

- ☐ Manufacturers of enterprise servers
- ☐ Manufacturers of entry servers
- ☐ Manufacturers of workstations
- ☒ Manufacturers of PC's
- ☐ Non-computer manufacturers
- ☐ Developers of operating systems
- ☒ Developers of networking software
- ☒ Developers of application software
- ☒ Integrated solution providers
- ☒ Service providers
- ☐ Other (Please specify below)

Patent Value Tool (Optional - this may be used by the inventor and attorney to assist with the evalua

(The Patent Value tool can be used by you or the evaluation team to determine the potential licensing value of your invention.)

The **Patent Value Tool** has not yet been used to calculate a score.

Post Disclosure Text & Drawings
Enter any additional information relating to this disclosure below:

(Form Revised 12/17/97)